nature physics

Article

https://doi.org/10.1038/s41567-026-03389-y

Simple and efficient end-to-end quantum thermal and ground state preparation

Received: 14 December 2025

Accepted: 25 June 2026

Published online: 12 August 2026

Check for updates

Zhiyan Ding [1,7] ✉, Yongtao Zhan [2,3,7], John Preskill [2,3,4] & Lin Lin [5,6]

Quantum computing algorithms for many-body physics, chemistry and materials science typically require the preparation of thermal or ground states for a given Hamiltonian. Here we propose quantum algorithms based on system–bath interactions to prepare thermal and ground states for a range of physically relevant Hamiltonians. These algorithms require only forward evolution under a system–bath Hamiltonian in which the bath is a single reusable ancilla qubit, making them especially well suited for early fault-tolerant quantum devices. By carefully designing the bath and interaction Hamiltonians, we prove that the fixed point of the dynamics accurately approximates the desired quantum state. Furthermore, we establish theoretical guarantees on the mixing time, and thereby provide a rigorous justification for the end-to-end efficiency of system–bath interaction models in thermal and ground state preparation for the physically relevant models we consider.

Preparing quantum thermal and ground states is a fundamental task in quantum computing, with wide-ranging applications in quantum many-body physics, quantum chemistry and materials science. A variety of approaches have been developed in the literature, including variational algorithms[1], adiabatic methods[2,3], matrix function-based methods using quantum phase estimation and quantum singular value transformation[4–6], tomography-based quantum imaginary time evolution[7] and dissipative state preparation schemes[8,9], to name a few.

While each approach has its own strengths and limitations, there has been a persistent trade-off between simplicity, efficiency, generality and rigour. Namely, we would like to design algorithms that are simple enough to be implemented on early fault-tolerant quantum devices. The runtime should scale at most polynomially with relevant parameters such as system size and precision. The polynomial runtime scaling is not a trivial requirement: preparing the ground state of a general local Hamiltonian is Quantum Merlin-Arthur-hard in the worst case[10], and even preparing classical Gibbs states can be non-deterministic polynomial-hard[11,12]. Thus, polynomial-time quantum algorithms cannot be expected for all Hamiltonians. Our goal should be to develop algorithms that are efficient for broad classes of physically relevant Hamiltonians, and all of these should come with rigorous end-to-end performance guarantees. So far, we are not aware of any approach that simultaneously satisfies all of these criteria.

Among existing approaches, dissipative algorithms appear to be the most promising candidates. They can avoid variational parameter tuning, state tomography procedures and normalization factors that can grow exponentially with system size. In recent years, a new wave of approaches based on dissipative dynamics, such as the Lindblad dynamics, have been proposed to prepare thermal and ground states of general Hamiltonians[13–27]. In a broad sense, dissipative algorithms mimic how nature prepares thermal states. The system of interest is coupled to an external bath that induces dissipation. Importantly, dissipation here does not mean that the system state is post-selected on the basis of specific measurement outcomes. This issue arises in direct implementations of imaginary-time evolution on quantum computers, where the probability of obtaining the desired thermal state in the system register decreases exponentially with system size. By contrast, dissipative algorithms are naturally described by quantum channels, ensuring that at each step the system register remains in a well-defined physical quantum state without post-selection. In this way,

[1]Department of Mathematics, University of Michigan, Ann Arbor, MI, USA. [2]Institute for Quantum Information and Matter, California Institute of Technology, Pasadena, CA, USA. [3]Division of Physics, Mathematics, and Astronomy, California Institute of Technology, Pasadena, CA, USA. [4]AWS Center for Quantum Computing, Pasadena, CA, USA. [5]Department of Mathematics, University of California, Berkeley, Berkeley, CA, USA. [6]Applied Mathematics and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. [7]These authors contributed equally: Zhiyan Ding, Yongtao Zhan. ✉e-mail: zyding@umich.edu

https://doi.org/10.1038/s41567-026-03389-y

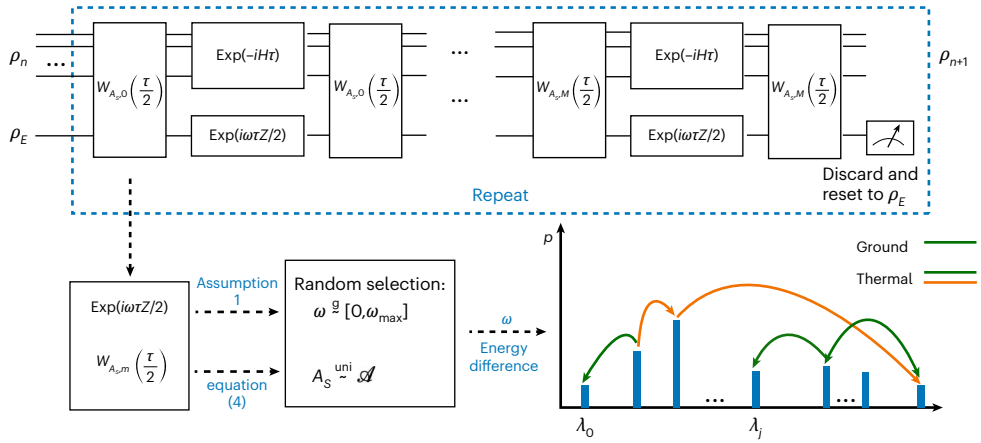

Fig. 1 | Illustration of the quantum circuit for the proposed algorithm. The weak-coupling dynamics in equation (2) are implemented via a second-order Trotter decomposition. The procedure requires only a single ancilla qubit. Here, $\rho_E \propto \exp(\beta\omega Z/2)$ for thermal state preparation, and $\rho_E = |0\rangle\langle 0|$ for ground state preparation. The system–bath interaction term $W_{A_S,m}(\tau/2)$ is defined in equation (5) with A_S randomly sampled from a set $\mathcal{A}$. For each $\omega > 0$ randomly sampled from a given distribution $g(\omega)$, the algorithm approximately generates a jump from an energy eigenstate $|\psi_j\rangle$ to another eigenstate $|\psi_i\rangle$ with Bohr frequency $\lambda_i - \lambda_j \approx \pm\omega$. That is, the change in energy of the system can be either $+\omega$ or $-\omega$. In the thermal state preparation algorithm, both energy-increasing and energy-decreasing transitions are allowed, while in the ground state preparation algorithm, only energy-decreasing transitions ($\lambda_i - \lambda_j \approx -\omega$) are permitted. The system–bath interaction block $W(\cdot)$ is explained in Methods.

dissipative algorithms can be interpreted as quantum generalizations of Markov chain Monte Carlo methods. By carefully designing the dissipation mechanism, or equivalently the quantum moves, and repeatedly applying the resulting quantum channel that simulates the dissipative dynamics, the system can be driven towards its thermal or ground state.

A key metric for evaluating the efficiency of dissipative algorithms is the total simulation time required to reach the target state. The total simulation time is often characterized by the mixing time, which refers to how long it takes for the system to evolve close enough to its steady state, regardless of the initial state. In other words, it measures the time needed for the system to 'forget' its starting point and settle into thermal or ground state equilibrium. Recently, there has been substantial progress in rigorous understanding of the mixing time of the continuous-time Lindblad dynamics for several classes of physically relevant Hamiltonians[28–34].

A main drawback of dissipative algorithms for thermal and ground state preparation lies in the complexity of their implementation, that is, the implementation of quantum channels that accurately simulate the dissipative dynamics. Existing high-order simulation algorithms[16,19,35–37] can achieve optimal or near-optimal scaling in terms of the simulation time, but they often involve complex logic gates, controlled or time-reversed Hamiltonian simulations, as well as a large number of ancilla qubits. These complex circuit structures can be challenging to implement on early fault-tolerant quantum devices, where quantum resources remain limited, error correction may be applied selectively to components most susceptible to noise, and hybrid architectures combining analogue (such as forward Hamiltonian evolution) and digital elements may be used to reduce implementation costs and suppress errors. Consequently, algorithms must be substantially simplified before they can be considered for practical deployment. Meanwhile, simplified protocols often deviate too far from the continuous-time Lindblad dynamics, making it difficult to rigorously analyse the mixing time.

In this work, we resolve the above tension by introducing a quantum algorithm that is simple to implement, and yet admits rigorous end-to-end performance guarantees. Our approach is a randomized system–bath interaction scheme that requires only forward evolution under the system Hamiltonian and a single reusable environment qubit. At a high level, the protocol proceeds as follows (Fig. 1). Given a system Hamiltonian H and a small set of local system operators $\mathcal{A}$, prepare a one-qubit environment state. In each round, draw a coupling $A_S \in \mathcal{A}$ and a random bath frequency ω from a prescribed distribution, evolve forward under a weak system–bath interaction with a smooth temporal envelope while the system simultaneously evolves under H, then trace out and reset the environment qubit. The resulting quantum channel acting on the system is iterated, and the number of iterations is controlled by the mixing time.

To justify the efficiency of this protocol, we compare this quantum channel with a continuous-time Lindblad evolution whose fixed point is known to approximate the target state. By carefully choosing the environment state, coupling operators and sampling distributions, we ensure that the fixed point of the channel can be made arbitrarily close to the target state. Moreover, we establish polynomial mixing-time bounds for several physically relevant Hamiltonians, thereby providing rigorous guarantees within a minimal architecture that is well-suited for early fault-tolerant devices.

We note that the idea of employing a weak system–bath interaction framework to cool a system and prepare low-temperature thermal or ground states has been explored in several previous and concurrent works[27,38–42], from both theoretical and numerical perspectives. In particular, refs. 39,41 adopted similar principles to design quantum algorithms for thermal state preparation based on system–bath interaction models, although some specific details of their protocols differ from ours. In this work, we carefully design both the initialization of the bath state and the interaction function, thereby providing an end-to-end efficiency guarantee within the system–bath interaction framework. Specifically, our study is the first to establish rigorous mixing-time guarantees for both thermal and ground state preparation in such a framework, laying a theoretical foundation for understanding the promising empirical results observed in related studies. A more detailed comparison with the previous work can be found in Supplementary Information section 2.

Algorithm description

Given a system Hamiltonian H, we define the total Hamiltonian as

$$H_\alpha(t) = H + H_E + \alpha f(t)\left(A_S \otimes B_E + A_S^\dagger \otimes B_E^\dagger\right). \tag{1}$$

Here, A_S is a coupling operator acting on the system, which can be randomly chosen from a user-provided set of operators $\mathcal{A} = \{A^i, -A^i\}_i$ with the property that $\{(A^i)^\dagger\}_i = \{A^i\}_i$. We assume without loss of gen-

erality $\|A^i\| \le 1$. We set $B_E = (X - iY)/2 = |1\rangle\langle 0|$. The bath Hamiltonian $H_E = -\omega Z/2$ depends on a random frequency variable ω, drawn from a distribution $g(\omega)$ related to the system Hamiltonian. Here, X, Y and Z denote the single-qubit Pauli operators. The interaction envelope $f(t)$ is a real, positive, even and normalized function of time, satisfying $f(t) = f(-t) \in \mathbb{R}$ and $\int_{\mathbb{R}} f^2(t)\,\mathrm{d}t = 1$. We choose f to be a Gaussian function $f(t) = \frac{1}{(2\pi)^{1/4}\sigma^{1/2}} \exp(-\frac{t^2}{4\sigma^2})$.

Starting from an initial state ρ_0 in the system register, in each step of the algorithm we set the state of the ancilla qubit to $\rho_E \propto \exp(-\beta H_E)$, simulate the time evolution governed by the total Hamiltonian $H_\alpha(t)$ and then trace out the ancilla qubit. When $\beta = \infty$, we set $\rho_E = |0\rangle\langle 0|$. Then,

$$\rho_{n+1} = \Phi(\rho_n) := \mathbb{E}_{A_S,\omega}\left(\mathrm{Tr}_E\left[U^\alpha(T)\,(\rho_n \otimes \rho_E)\,U^\alpha(T)^\dagger\right]\right), \qquad (2)$$

where $U^\alpha(t)$ is the time-evolution operator of the total Hamiltonian $U^\alpha(t) := \mathcal{T}\exp(-i\int_{-T}^{t} H_\alpha(s)\,\mathrm{d}s)$. Note that the Hamiltonian evolution starts at time $-T$. In the implementation, we simulate the time-dependent Hamiltonian in Φ using exisiting Hamiltonian simulation protocols. A detailed discussion can be found in Methods and Supplementary Information section 6D.

This work provides a rigorous analysis showing that, for physically relevant but analytically tractable classes of models such as free fermions and commuting Hamiltonians, the map Φ possesses a unique fixed point that accurately approximates the desired thermal or ground state, and that its mixing time can be explicitly bounded. These results offer a concrete demonstration that the mechanism underlying Φ can be made fully rigorous in non-trivial quantum many-body settings. More generally, we prove that, whenever the mixing time τ_{mix} of Φ grows at most polynomially in the system size, one can choose parameters so that the state $\rho_{\tau_{\mathrm{mix}}}$ stays arbitrarily close to the fixed point ρ_{fix} and, hence, serves as a reliable approximation to the target thermal or ground state.

Rigorous guarantees

Our first result is to show that the fixed point of the quantum channel Φ can be made arbitrarily close to the target state. The mixing time $\tau_{\mathrm{mix}}(\epsilon)$ is defined as the minimum number of iterations required to ensure that $\Phi^k(\rho_0)$ is ϵ−close to the fixed point for any initial state ρ_0 (see Definition S1 in Supplementary Information for the formal definition).

Theorem 1. (Informal main result) *Let Φ be the quantum channel in equation (2). For any target precision $\epsilon > 0$, there exist parameter choices such that the trace distance between the fixed point $\rho_{\mathrm{fix}}(\Phi)$ and the target state is at most ϵ. Starting from an arbitrary initial state, the end-to-end Hamiltonian evolution time is upper bounded by*

$$T_{\mathrm{total}} = \begin{cases} \mathrm{poly}\left(\tau_{\mathrm{mix}}(\epsilon), \beta, 1/\epsilon\right), & \text{thermal}, \\ \mathrm{poly}\left(\tau_{\mathrm{mix}}(\epsilon), \Delta^{-1}, 1/\epsilon\right), & \text{ground state}, \end{cases} \qquad (3)$$

where β is the inverse temperature and Δ is the spectral gap of the system Hamiltonian H.

We note that ground states can also be approximated by thermal states at low enough temperatures, $\beta = \mathrm{poly}(N, 1/\Delta)$ (ref. 43). We provide a detailed discussion of this complexity discrepancy in Supplementary Information section 5. The total Hamiltonian evolution time T_{total} in theorem 1 is the product of the number of iterations $\tau_{\mathrm{mix}}(\epsilon)$ and the Hamiltonian evolution time $2T$ in each iteration. The polynomial dependence on the mixing time $\tau_{\mathrm{mix}}(\epsilon)$ occurs because the Hamiltonian evolution time T in each iteration needs to scale polynomially with $\tau_{\mathrm{mix}}(\epsilon)$, so that the fixed point of Φ can remain close to the target state after $\tau_{\mathrm{mix}}(\epsilon)$ iterations.

To establish end-to-end efficiency in preparing thermal and ground states, we derive rigorous mixing-time bounds for a range of physically relevant Hamiltonians. These results are presented in the following two theorems:

Theorem 2. *For thermal state preparation of N-qubit non-interacting spin systems or quadratic fermionic systems with $\beta = \Theta(1)$, and of local commuting spin systems at high temperature, the mixing* time *$\tau_{\mathrm{mix}}(\epsilon)$* of *$\Phi$ scales as poly$(N, 1/\epsilon)$.*

Theorem 3. *For ground state preparation of N-qubit non-interacting spin systems or quadratic fermionic systems, the mixing time $\tau_{\mathrm{mix}}(\epsilon)$ of Φ scales as poly$(N, 1/\epsilon)$.*

By combining the mixing-time bounds with theorem 1, the algorithm achieves runtime scaling polynomially with system size.

Analysis overview

We first need to show that the fixed point of Φ is ϵ-close to the target thermal or ground state, and then derive an upper bound on the mixing time of Φ.

First, we show that equation (2) can be approximated by a continuous-time Lindbladian dynamics whose fixed point can be tuned to lie arbitrarily close to the target state by setting $T = \tilde{\Omega}(\sigma)$ and $\sigma \gg 1$, $\alpha T \ll 1$. Specifically, we first show that Φ can be approximated by a Lindbladian dynamics $\exp(\mathcal{L}\alpha^2)$–up to a unitary evolution—with error $\mathcal{O}(\alpha^4\sigma^2)$, where $\mathcal{L}$ is the Lindbladian operator and α^2 corresponds to the Lindbladian evolution time. This result is formally given in Supplementary Information section 3. Next, we show that the fixed point of the Lindbladian dynamics closely approximates the target thermal or ground state. Consequently, the fixed point of Φ is also close to the desired state. For the thermal state, the fixed-point error can be bounded by $\widetilde{\mathcal{O}}((\beta/\sigma)\alpha^2\tau_{\mathrm{mix}}(\epsilon))$, while for the ground state, the error scales as $\mathcal{O}(\exp(-\sigma^2\Delta^2/32)\alpha^2\tau_{\mathrm{mix}}(\epsilon))$, provided the Hamiltonian has a spectral gap Δ. The detailed versions of the approximation error are stated in Supplementary Information section 5 as theorem S4 and theorem S5, respectively.

According to the approximation error bound, when $\alpha^2\tau_{\mathrm{mix}}(\epsilon)$ remains essentially constant, an appropriate choice of the parameters σ and α ensures that the fixed point $\rho_{\mathrm{fix}}(\Phi)$ can be made arbitrarily close to the target state. Moreover, because Φ can be approximated by the Lindbladian dynamics, we show that the rescaled mixing time $\alpha^2\tau_{\mathrm{mix}}(\epsilon)$ closely matches the mixing time of the Lindbladian operator $\mathcal{L}$. This correspondence is a key step in establishing an upper bound on the mixing time of Φ in the second stage of our analysis.

The main technical challenge, however, is that in practice the rescaled mixing time $\alpha^2\tau_{\mathrm{mix}}(\epsilon)$ may diverge as σ or α^{-1} grows. For instance, in thermal state preparation the fixed-point error bound often scales as $\alpha^2\tau_{\mathrm{mix}}(\epsilon)/\sigma$. Without a carefully designed filter function f and system–bath interaction, $\alpha^2\tau_{\mathrm{mix}}(\epsilon)$ can grow exponentially with σ, even for the simple single-qubit Hamiltonian $H = Z$ (Supplementary Information section 7, theorem S17). Thus a key step in our work is to design the dissipative protocol to ensure that once σ is sufficiently large, the mixing time $\tau_{\mathrm{mix}}(\epsilon)$ is independent of σ.

To achieve this, we need to carefully design the initial bath state and the interaction function. This construction provides substantial flexibility in tuning f, allowing the resulting Lamb-shift term to approximately commute with the target thermal or ground state. Specifically, for a fixed ω and sufficiently large σ, the system–bath interaction is designed so that the induced energy transitions in the system concentrate around ω rather than near zero. Then, by sampling ω from a suitable distribution, our method enables large energy transitions while keeping the perturbation to the approximate fixed point small, thereby ensuring provably fast mixing without substantially disturbing the target state.

Applications of the algorithm

Quantum materials and chemistry on early fault-tolerant hardware

Efficient state preparation with end-to-end performance guarantees is essential for reliable property prediction in materials and molecules,

especially when classical methods disagree. Our protocol prepares thermal and ground states using only forward evolution under a system–bath Hamiltonian and a single ancilla qubit. This minimal architecture reduces control overhead and avoids deep control logic, which is attractive for early fault-tolerant and hybrid analogue–digital devices.

Benchmarks for dissipative engineering and codesign with error correction

Because the channel's fixed point can be tuned to approximate a target thermal or ground state with minimal circuitry, our protocol provides a benchmark that directly reflects the end-to-end performance of a quantum device. This perspective is especially relevant for architectures with mid-circuit measurement capabilities, which is a core ingredient of error correction. Rather than focusing on per-gate fidelity, the dissipative dynamics test how well the hardware and control stack jointly support long-time evolution towards a stable fixed point. In practice, the most natural initial targets are Hamiltonians with simple native structure, such as free-fermion or commuting local models, and modest system sizes where forward system–bath evolution and repeated ancilla reset can be implemented with minimal ancilla counts and low overhead. In this sense, the protocol is a natural setting for codesign: it allows hardware and algorithmic choices to be evaluated together in terms of their impact on the overall ability to prepare and stabilize complex quantum states.

Conclusion and outlook

In this work, we have presented a quantum algorithm for preparing both thermal and ground states with provable correctness and mixing-time guarantees, while requiring only modest quantum resources. The algorithm operates by repeatedly applying a simple quantum channel, obtained from forward time evolution under a system–bath Hamiltonian, followed by resetting the bath. The bath can be reduced to a single ancilla qubit by employing a randomized choice of bath frequency and coupling operator between the system and bath. This design makes it well suited for early fault-tolerant quantum devices, including hybrid devices that can implement certain type of Hamiltonian simulations via analogue methods.

Our analysis shows that, with a suitable choice of parameters, the fixed point of the quantum channel can accurately approximate the target thermal or ground state. Once the fixed-point error bound is established, it suffices to prove polynomial mixing time in order to achieve end-to-end efficiency. A key technical point is that the mixing time and simulation parameters are generally interdependent; it is therefore important to demonstrate that, for our channel construction, the mixing time remains bounded as parameters are tuned to reduce the fixed-point error. While we have established polynomial mixing for several physically relevant models, similar behaviour may also extend to broader classes of Hamiltonians, especially in the thermal-state setting, as indicated by two recent works[44,45] that appeared after the initial completion of this paper.

The weak system–bath interaction framework can be regarded as an approximate, low-order approach to Lindblad simulation. In this work, we propose a second-order Trotterization method for the simulation of Φ, which circumvents the need for backward Hamiltonian evolution. Compared with high-order Lindblad simulation methods[35–37], this can lead to an asymptotic slowdown with respect to parameters such as system size and precision. This tradeoff is analogous to the use of low-order Trotter methods in Hamiltonian simulation, and is to some extent unavoidable. Nevertheless, there is considerable potential for improvement to reduce the cost. For instance, if backward Hamiltonian evolution can be efficiently implemented on the device, one may employ a higher-order Trotterization scheme to simulate Φ, thereby substantially reducing the asymptotic gate count. This refinement further enhances the algorithms practicality for early fault-tolerant implementations across diverse quantum hardware platforms.

In the hardware implementation, conditioned on sampling a pure eigenstate from ρ_E and on a measurement outcome for the discarded ancilla, each iteration of our algorithm maps a pure system state to a pure system state. This process naturally induces a classical Markov chain on the space of state vectors, with each application of Φ corresponding to a single transition step. Consequently, this protocol offers a quantum-inspired approach to sampling from Gibbs ensembles using classical state-vector or tensor-network simulations, in regimes where both the states and transitions admit efficient representations. This can also lead to a hybrid workflow: using classical resources to pre-screen couplings and frequency windows, followed by quantum preparation at scales where classical methods saturate.

In our implementation, the filter function has a narrow width σ^{-1} in the frequency domain. This choice simplifies the fixed-point error analysis but may increase the total simulation time. Allowing broader energy transitions per channel application, as in the protocols of [18,19] for ground- and thermal-state preparation, could lead to a more efficient implementation. However, this also introduces new challenges, such as demonstrating that the approximating Lindbladian dynamics still fixes the target state. We hope this work will motivate further research on optimized filter design, generalized mixing-time analysis and experimental demonstrations of dissipative algorithms for quantum state preparation.

Online content

Any methods, additional references, Nature Portfolio reporting summaries, source data, extended data, supplementary information, acknowledgements, peer review information; details of author contributions and competing interests; and statements of data and code availability are available at https://doi.org/10.1038/s41567-026-03389-y.

References

1. Tilly, J. et al. The Variational Quantum Eigensolver: a review of methods and best practices. *Phys. Rep.* https://doi.org/10.1016/j.physrep.2022.08.003 (2022).
2. Farhi, E., Goldstone, J., Gutmann, S. & Sipser, M. Quantum computation by adiabatic evolution. Preprint at https://arxiv.org/abs/quant-ph/0001106 (2000).
3. Albash, T. & Lidar, D. A. Adiabatic quantum computation. *Rev. Mod. Phys.* **90**, 015002 (2018).
4. Poulin, D. & Wocjan, P. Preparing ground states of quantum many-body systems on a quantum computer. *Phys. Rev. Lett.* **102**, 130503 (2009).
5. Gilyén, A., Su, Y., Low, G. H. & Wiebe, N. Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics. In *Proc. 51st Annual ACM SIGACT Symposium on Theory of Computing* 193–204 (Association for Computing Machinery, 2019).
6. Lin, L. & Tong, Y. Near-optimal ground state preparation. *Quantum* **4**, 372 (2020).
7. Motta, M. Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution. *Nat. Phys.* **16**, 205–210 (2020).
8. Verstraete, F., Wolf, M. M. & Cirac, I. Quantum computation and quantum-state engineering driven by dissipation. *Nat. Phys.* **5**, 633–636 (2009).
9. Temme, K., Osborne, T. J., Vollbrecht, K. G., Poulin, D. & Verstraete, F. Quantum Metropolis sampling. *Nature* **471**, 87–90 (2011).
10. Kitaev, A. Y., Shen, A. & Vyalyi, M. N. *Classical and Quantum Computation* (American Mathematical Society, 2002).
11. Barahona, F. On the computational complexity of Ising spin glass models. *J. Phys. A* **15**, 3241 (1982).
12. Sly, A. Computational transition at the uniqueness threshold. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science* 287–296 (IEEE Computer Society, 2010).

13. Cubitt, T. S. Dissipative ground state preparation and the dissipative quantum eigensolver. Preprint at https://arxiv.org/abs/2303.11962 (2023).
14. Mi, X. Stable quantum-correlated many-body states through engineered dissipation. *Science* **383**, 1332–1337 (2024).
15. Rall, P., Wang, C. & Wocjan, P. Thermal state preparation via rounding promises. *Quantum* **7**, 1132 (2023).
16. Chen, C.-F., Kastoryano, M. J., Brandão, F. G. & Gilyén, A. Quantum thermal state preparation. Preprint at https://arxiv.org/abs/2303.18224 (2023).
17. Chen, C.-F., Kastoryano, M. J. & Gilyén, A. An efficient and exact noncommutative quantum Gibbs sampler. Preprint at https://arxiv.org/abs/2311.09207 (2023).
18. Ding, Z., Li, B. & Lin, L. Efficient quantum Gibbs samplers with Kubo–Martin–Schwinger detailed balance condition. *Commun. Math. Phys.* **406**, 67 (2025).
19. Ding, Z., Chen, C.-F. & Lin, L. Single-ancilla ground state preparation via Lindbladians. *Phys. Rev. Research* **6**, 033147 (2024).
20. Li, H.-E., Zhan, Y. & Lin, L. Dissipative ground state preparation in ab initio electronic structure theory. *npj Quantum Info.* **11**, 183 (2025).
21. Zhan, Y., et al. Rapid quantum ground state preparation via dissipative dynamics. *Physical Review X* **16**, 011004 (2026).
22. Hahn, D., Sweke, R., Deshpande, A. & Shtanko, O. Efficient quantum Gibbs sampling with local circuits. *PRX Quantum* **7**, 020314 (2026).
23. Gilyén, A., Chen, C.-F., Doriguello, J. F. & Kastoryano, M. J. Quantum generalizations of Glauber and Metropolis dynamics. Preprint at https://arxiv.org/abs/2405.20322 (2024).
24. Jiang, J. & Irani, S. Quantum Metropolis sampling via weak measurement. Preprint at https://arxiv.org/abs/2406.16023 (2024).
25. Eder, P. J., Finžgar, J. R., Braun, S. & Mendl, C. B. Quantum dissipative search via Lindbladians. *Phys. Rev. A* **111**, 042430 (2025).
26. Motlagh, D., Zini, M. S., Arrazola, J. M. & Wiebe, N. Ground state preparation via dynamical cooling. Preprint at https://arxiv.org/abs/2404.05810 (2024).
27. Lloyd, J., Michailidis, A. A., Mi, X., Smelyanskiy, V. & Abanin, D. A. Quasiparticle cooling algorithms for quantum many-body state preparation. *PRX Quantum* **6**, 010361 (2025).
28. Bardet, I. Rapid thermalization of spin chain commuting Hamiltonians. *Phys. Rev. Lett.* **130**, 060401 (2023).
29. Rouzé, C., França, D. S. & Alhambra, Á. M. Efficient thermalization and universal quantum computing with quantum gibbs samplers. In *STOC 25* 1488–1495 https://doi.org/10.1145/3717823.3718268 (Association for Computing Machinery, 2025);
30. Ding, Z., Landau, Z., Li, B., Lin, L. & Zhang, R. Polynomial-time preparation of low-temperature Gibbs states for 2D Toric Code. *J. Math. Physics* **67**, 052201 (2026).
31. Kochanowski, J., Alhambra, A. M., Capel, A. & Rouzé, C. Rapid thermalization of dissipative many-body dynamics of commuting Hamiltonians. *Commun. Math. Phys.* **407**, 176 (2024).
32. Rouzé, C., França, D. S. & Alhambra, Á. M. Optimal quantum algorithm for Gibbs state preparation. *Phys. Rev. Lett.* **136**, 060601 (2026).
33. Tong, Y. & Zhan, Y. Fast mixing of weakly interacting fermionic systems at any temperature. *PRX Quantum* **6**, 030301 (2025).
34. Šmíd, Š., Meister, R., Berta, M. & Bondesan, R. Polynomial-time quantum Gibbs sampling for the weak and strong coupling regime of the Fermi-Hubbard model at any temperature. *Nat. Comms.* **16**, 10736 (2025).
35. Cleve, R. & Wang, C. Efficient quantum algorithms for simulating Lindblad evolution. In *ICALP 2017* (eds Chatzigiannakis, I. et al.) Vol. 80, 17:1–17:14 (Schloss Dagstuhl Leibniz-Zentrum Informatik, 2017).
36. Li, X. & Wang, C. Simulating Markovian open quantum systems using higher-order series expansion. In *ICALP 2023* (eds Etessami, K., et al.) Vol. 261, 87:1–87:20 (Schloss Dagstuhl Leibniz-Zentrum Informatik, 2023).
37. Ding, Z., Li, X. & Lin, L. Simulating open quantum systems using Hamiltonian simulations. *PRX Quantum* **5**, 020332 (2024).
38. Molpeceres, D., Lu, S., Cirac, J. I. & Kraus, B. Quantum algorithms for cooling: a simple case study. *Phys. Rev. Res.* **7**, 033162 (2025).
39. Hahn, D., Parameswaran, S. A. & Placke, B. Efficient quantum thermal state preparation via local driving. *Phys. Rev. B* **114**, 014302 (2026).
40. Langbehn, J. et al. Universal cooling of quantum systems via randomized measurements. Preprint at https://arxiv.org/abs/2506.11964 (2025).
41. Lloyd, J. & Abanin, D. A. Quantum thermal state preparation for near-term quantum processors. Preprint at https://arxiv.org/abs/2506.21318 (2025).
42. Scandi, M. & M. Alhambra, Á. Thermalization in open many-body systems and KMS detailed balance. *Phys. Rev. X* **16**, 011040 (2026).
43. Kuwahara, T., Alhambra, ÁM. & Anshu, A. Improved thermal area law and quasilinear time algorithm for quantum gibbs states. *Phys. Rev. X* **11**, 011047 (2021).
44. Slezak, S., Scandi, M., M. Alhambra, Á., França, D. S. & Rouzé, C. Polynomial-time thermalization and Gibbs sampling from system–bath couplings. Preprint at https://arxiv.org/abs/2601.16154 (2026).
45. Wang, K. & Ding, Z. Beyond Lindblad dynamics: rigorous guarantees for thermal and ground state preservation under system bath interactions. Preprint at https://arxiv.org/abs/2512.03457 (2026).

Methods

Here, we first present the implementation details of the proposed quantum algorithm and then summarize the proof roadmap and the core analytical techniques supporting the main theoretical result.

Algorithm implementation

To facilitate early fault-tolerant implementation, we approximate the time-dependent Hamiltonian in Φ using a second-order Trotter formula with a small time step τ. We choose not to employ higher-order Trotter formulas so as to avoid implementing backward Hamiltonian simulations of H. At each time step n, the frequency parameter ω and the operator A_S are independently sampled. Let $M=\lceil 2T/\tau\rceil$, and $f_m=f((m+1/2)\tau-T)$. According to the preceding description, the evolution from ρ_n to ρ_{n+1} is implemented as follows:

$$\begin{aligned}\rho_{n+1} &= \Phi^{\text{approx}}(\rho_n)\\ &:= \mathbb{E}_{A_S,\omega}\left(\mathrm{Tr}_E\left[\Pi_{m=0}^{M} U_{m\tau}^{\alpha}(\rho_n\otimes\rho_E)\left(\Pi_{m=0}^{M}U_{m\tau}^{\alpha}\right)^{\dagger}\right]\right),\end{aligned} \tag{4}$$

where

$$U_{m\tau}^{\alpha} = W_{A_S,m}(\tau/2)\exp(-iH\tau)\exp(i\omega\tau Z/2)W_{A_S,m}(\tau/2)$$

implements the Trotterized approximation to U^{α} with

$$W_{A_S,m}(\tau/2) = \exp(-i\alpha f_m(A_S\otimes B_E + A_S^{\dagger}\otimes B_E^{\dagger})\tau/2). \tag{5}$$

In practice, if the system Hamiltonian H consists of noncommuting terms, one would further Trotterize $\exp(-iH\tau)$. When first- or second-order Trotterization is used for the Hamiltonian evolution, the simulation involves only forward time evolution, which is well suited to analogue quantum simulators. When backward Hamiltonian evolution can also be implemented efficiently, as in digital quantum computers, the simulation cost can be further reduced by using higher-order Trotterization or quantum singular value transformation methods to simulate Φ with improved precision. The detailed complexity analysis of the second-order Trotter can be found in Supplementary Information section 6D.

Key techniques

The analysis of our method can be divided into three main steps: (1) derive the effective Lindblad dynamics that approximates Φ; (2) show that the fixed point of the Lindblad dynamics is close to the target thermal (or ground) state and, hence, also serves as the fixed point of Φ; and (3) establish an upper bound on the mixing time of Φ.

Derivation of effective Lindblad dynamics. We show that the quantum map Φ in the weak coupling regime $\alpha\ll 1$ can be approximated by an effective Lindblad dynamics up to a unitary transformation.

Define the time evolution operator by $U_S(t) := \exp(-iHt)$, and the associated superoperator by $\mathcal{U}_S(t)[\rho] = U_S(t)\rho U_S^{\dagger}(t)$. Given a jump operator V, we define the associated dissipative operator as

$$\mathcal{D}_V(\rho) = V\rho V^{\dagger} - \frac{1}{2}\{V^{\dagger}V,\rho\}. \tag{6}$$

The following theorem establishes that the quantum channel Φ can be approximated by an effective Lindblad dynamics:

Theorem 4. (Informal) *Under the choice of $H_E, A_S, B_E, f(t), g(\omega)$ in the main text, ρ_{n+1} can be expressed as*

$$\rho_{n+1} = \mathcal{U}_S(T)\circ\exp(\mathcal{L}\alpha^2)\circ\mathcal{U}_S(T)[\rho_n] + \mathcal{O}(\alpha^4T^4\,\|f\|_{L^\infty}^4). \tag{7}$$

Here, $\mathcal{L}$ is a Lindbladian operator.

$$\begin{aligned}\mathcal{L}(\rho) = \mathbb{E}_{A_S}\Bigg(\int_{-\infty}^{\infty} -i\left[g(\omega)H_{\mathrm{LS},A_S}(\omega),\rho\right]\\ +\gamma(\omega)\mathcal{D}_{V_{A_S f,T}(\omega)}(\rho)\mathrm{d}\omega\Bigg).\end{aligned} \tag{8}$$

Here, the Lamb shift term $H_{\mathrm{LS},A_S}(\omega)$ is a Hermitian matrix. $\gamma(\omega) = (g(\omega)+g(-\omega))/(1+\exp(\beta\omega))$ when $\beta<\infty$, and $\gamma(\omega)=(g(\omega)+g(-\omega))\mathbf{1}_{\omega<0}+g(0)\mathbf{1}_{\omega=0}$ when $\beta=\infty$. The jump operator $V_{A_S f,T}(\omega)$ is defined as

$$V_{A_S f,T}(\omega) = \int_{-T}^{T} f(t)A_S(t)\exp(-i\omega t)\mathrm{d}t, \tag{9}$$

where $A_S(t)$ is the Heisenberg-picture evolution of the coupling operator A_S: $A_S(t) = \exp(iHt)A_S\exp(-iHt)$.

In the above theorem, we omit the f, T dependence of $H_{\mathrm{LS},A_S}(\omega)$ in the subindex for simplicity. The rigorous version of above theorem with the specific form of $\mathcal{L}$ can be found in Supplementary Information section 3 as theorem S2. The $\mathcal{O}(\alpha^2)$ term is obtained by analysing the second-order term in the Dyson expansion of the time evolution operators, and the $\mathcal{O}(\alpha^4)$ remainder bounds the sum of higher-order terms. The parameter α^2 can be interpreted as the effective evolution time under the Lindbladian operator $\mathcal{L}$. Each application of Φ approximately evolves the system, up to the unitary transformation $\mathcal{U}_S(T)$, under the Lindbladian operator $\mathcal{L}$ for time α^2 with error scaling as $\mathcal{O}(\alpha^4)$. Therefore, up to the unitary transformation $\mathcal{U}_S(T)$, the map Φ serves as a first-order approximation to the Lindblad dynamics $\partial_t\rho=\mathcal{L}(\rho)$.

Choice of interaction terms. In our algorithm, it is important to properly choose $H_E, \rho_E, A_S, B_E, \alpha, f(t)$ and $g(\omega)$ to ensure that the fixed point of Φ closely approximates the target thermal or ground state and that the resulting dynamics has a favourable mixing time. First, H_E specifies the single-ancilla bath system. The operators A_S and B_E jointly induce energy transitions through the system–bath interaction, with the transition energy set by the parameter ω in H_E, thereby driving transitions between eigenstates of the system Hamiltonian. The bath is initialized in the thermal state ρ_E of H_E, which ensures the correct balance between upward and downward transitions, as reflected by $\gamma(\omega)$ in equation (8), and thus yields an approximate quantum detailed balance condition. The choice of coupling operator A_S determines which transitions between system eigenvectors are enabled. In particular, if all operators in $\mathcal{A}$ commute with H, then the interaction cannot induce transitions between distinct energy eigenspaces, so such a choice is not suitable for thermalization or cooling. More generally, for the corresponding effective Lindbladian with a full-rank fixed point, a standard sufficient condition for uniqueness is that the commutant of the coherent term together with the jump operators is trivial[23,46,47]. When prior knowledge of the eigenstate structure is available, one can design A_S to optimally induce transitions between relevant eigenstates. More generally, the Lindbladian literature provides generic choices of A_S that yield efficient local transitions for broad classes of systems, such as $\mathcal{A}=\{\pm X_i,\pm Y_i,\pm Z_i\}$ for local spin systems[48,49] and $\mathcal{A}=\{\pm c_i,\pm c_i^{\dagger}\}$ for fermionic systems[33,49]. The Fourier transform $\hat{f}$ of the filter function controls the allowed energy-transition window through the jump operator $V_{A_S f,T}(\omega)$ in equation (9). As shown in proposition 5, to guarantee that the fixed point is close to the target state, we need to choose σ sufficiently large, which generates a narrow energy-transition window around ω.

Therefore, to keep the mixing time bounded as σ increases, we must carefully design the system–bath interaction so that the induced energy transitions can cover a wide energy range; this is ensured by the choice of $g(\omega)$. The choice of $g(\omega)$ is flexible, as long as it provides a sufficiently broad energy-transition window. This choice does not require computing the spectrum of H. For a local Hamiltonian $H=\sum_j h_j$ with $\|h_j\|=\mathcal{O}(1)$, one may simply take $\omega_{\max}=2\sum_j\|h_j\|=\mathcal{O}(N)$, which upper bounds all Bohr frequencies. When A_S induces only local transitions, we expect that it is sufficient to choose $g(\omega)$ to be uniform on a smaller window with $\omega_{\max}=\mathcal{O}(1)$ to cover the local energy changes generated by A_S. In Supplementary Information section 7, we use a single qubit example to explicitly demonstrate the importance of this design choice.

Approximate fixed point. For a quantum channel Φ with a unique fixed point ρ_{fix}, a state ρ is close to ρ_{fix} if $\Phi(\rho) \approx \rho$. To ensure that ρ_{fix} approximates the target thermal state ρ_β, it suffices to bound $\|\Phi(\rho_\beta) - \rho_\beta\|_1$. According to theorem 4 and the fact that $\mathcal{U}_S(T)$ preserves the targe state, we can show

$$\| \Phi(\rho_\beta) - \rho_\beta \|_1 \approx \alpha^2 \left\| \mathcal{L}(\rho_\beta) \right\|_1, \quad \alpha \ll 1,$$

where $\mathcal{L}$ is the approximation Lindblad operator that appears in theorem 4. Thus, it suffices to show that the Lindblad dynamics approximately preserve the thermal or ground state. This forms the most technical component of the proof. For the thermal state, we demonstrate that the Lindbladian operator approximately satisfies the detailed balance condition, up to an additional coherent term that nearly commutes with the thermal state. For the ground state, assuming the Hamiltonian has a spectral gap Δ, we directly show that the operator approximately preserves the ground state.

The following proposition provides fixed point error bounds in preparing both thermal and ground states. Guidance on parameter choices and a more detailed formulation of this proposition can be found in theorem S4 (theorem S7) and theorem S5 (theorem S11).

Proposition 5. For the quantum channel defined in equation (2), For any inverse temperature $\beta > 0$, we have

$$\| \rho_{\text{fix}}(\Phi) - \rho_\beta \|_1 = \widetilde{\mathcal{O}}\left(\left(\frac{\beta}{\sigma} + \alpha^2 \sigma^2 \right) \alpha^2 \tau_{\text{mix}}(\epsilon) \right). \tag{10}$$

For ground state that $\beta = \infty$, if the Hamiltonian H has spectral gap Δ, we also have

$$\| \rho_{\text{fix}}(\Phi) - \rho_\infty \|_1 = \widetilde{\mathcal{O}}\left(\left(\mathrm{e}^{-\sigma^2 \Delta^2/32} + \alpha^2 \sigma^2 \right) \alpha^2 \tau_{\text{mix}}(\epsilon) \right). \tag{11}$$

Mixing time bound. The ground state and thermal state cases require different proof strategies.

We start with thermal state case. When two quantum channels Φ_1 and Φ are close, that is, $\|\Phi_1 - \Phi\|_{1\to 1} \approx 0$, their mixing times $\tau_{1,\text{mix}}(\epsilon), \tau_{\text{mix}}(\epsilon)$ are also close. We therefore define the quantum channel Φ_1

$$\Phi_1 = \mathcal{U}_S(T) \circ \exp\left(\mathcal{M}\alpha^2\right) \circ \mathcal{U}_S(T),$$

where $\mathcal{M}$ contains a dissipative operator that approximately satisfies the quantum detailed balance condition, together with a coherent term that commutes with the thermal state. In the free-fermionic example, we show that the dissipative part approximately satisfies the KMS detailed balance condition and possesses a non-vanishing spectral gap of order poly$(1/N)$ as $\sigma \to +\infty$. For commuting local Hamiltonians, existing results have established that the Davies generator exhibits a bounded mixing time in the high-temperature regime. We note that, even though the dissipative part possesses a spectral gap, existing techniques for analysing the mixing time cannot be directly applied due to the presence of the coherent term and the accompanying unitary evolution. To address this challenge, we establish the contractivity of the channel under a novel weighted Hilbert–Schmidt norm, $\| \rho_\beta^{-1/4}[\cdot]\rho_\beta^{-1/4} \|_2$, inspired by ref. 16. We prove contraction of the quantum channel $\exp(\mathcal{M}\alpha^2)$ under the weighted Hilbert–Schmidt norm, and this contraction still holds in the presence of the unitary evolution $\mathcal{U}_S(T)$ and therefore also holds for the map Φ_1. The Lindbladian $\mathcal{M}$ is chosen to both preserve the thermal state ρ_β and approximate the effective Lindbladian $\mathcal{L}$. Hence, the mixing time of Φ_1 is close to that of Φ. Using this framework, we rigorously demonstrate that, at high temperature, the mixing time for the single-qubit case, free-fermion systems, and commuting local Hamiltonians scales polynomially with N and remains independent of σ. The detailed results are summarized in Supplementary Information section 6.

For the ground state case, we focus only on the free fermion Hamiltonian and adopt the strategy from ref. 21 (section IV), which analyses the Heisenberg evolution of the number operator. Following the argument in ref. 21 (section IV), the convergence of the Lindblad dynamics to the ground state can be established by showing that the expectation of the number operator converges to zero. Moreover, because the unitary evolution commutes with the number operator, it does not affect this convergence. Finally, the convergence of the number operator can be directly related to the trace distance between the current state and the ground state using the Fuchs–van de Graaf inequality (see Supplementary Information section 8 for details).

Data availability

No datasets were generated or analysed during the current study.

References

46. Wolf, M. M. Quantum channels & operations: guided tour. https://mediatum.ub.tum.de/1701036 (Niels-Bohr Institute, 2012).
47. Zhang, Y. & Barthel, T. Criteria for Davies irreducibility of Markovian quantum dynamics. *J. Phys. A* **57**, 115301 (2024).
48. Temme, K. & Kastoryano, M. J. How fast do stabilizer Hamiltonians thermalize? Preprint at https://arxiv.org/abs/1505.07811 (2015).
49. Temme, K., Pastawski, F. & Kastoryano, M. J. Hypercontractivity of quasi-free quantum semigroups. *J. Phys. A* **47**, 405303 (2014).

Acknowledgements

We thank A. Chen, D. Hahn, J. Jiang, J. Lloyd, S. Parameswaran and O. Shtanko for helpful discussions. The Institute for Quantum Information and Matter is an NSF Physics Frontiers Center. L.L. is a Simons Investigator in Mathematics.

Author contributions

Z.D. and Y.Z. developed the original algorithm. Z.D., Y.Z. and L.L. carried out the theoretical analysis. J.P. and L.L. supervised the project. All authors contributed to the discussion and interpretation of the results and to the preparation and revision of the manuscript.

Funding

This material is based upon work supported by the US Department of Energy, Office of Science, Accelerated Research in Quantum Computing Centers, Quantum Utility through Advanced Computational Quantum Algorithms, grant no. DE-SC0025572 (J.P. and L.L.) and Fundamental Algorithmic Research toward Quantum Utility, grant no. DE-SC0025535 (J.P.). Additional support is from the US Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Systems Accelerator (Z.D., Y.Z., J.P. and L.L.) and the National Science Foundation, grant no. PHY-2317110 (Y.Z. and J.P.).

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at https://doi.org/10.1038/s41567-026-03389-y.

Correspondence and requests for materials should be addressed to Zhiyan Ding.

Peer review information *Nature Physics* thanks Xinzhao Wang and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.